

REMARKS

Claims 1, 4, 6-13, 15-18 and 21-26 are pending in the patent application.

The Examiner had newly rejected Claims 1, 4, 6-13, 15-18 and 21-26 under 35 USC 103(a) as unpatentable over Yang in view of Chen and further in view of Miller.

A telephone interview was conducted on August 19, 2008 to discuss a "Proposed Amendment". Applicants thank the Examiner for his time and consideration of the proposed amendments to the claims. Applicants herein formally submit those amendments and respectfully conclude that the claims, as amended, are patentable over the cited art.

The present invention is directed to an apparatus, program storage device, and a method for evaluating workload across a processing environment having a plurality of computer systems each having a plurality of assigned workload units, wherein the method comprises the steps of assigning a plurality of impact values, one impact value for each workload unit assigned for each of the plurality of computing systems, wherein the assigning of each impact value comprises determining a change in system expiration date should the workload unit be removed from the system; and assessing the workload based on the impact values. An impact value is assigned for each workload unit, wherein a

Y0R920000461**US1** **-16-**

workload unit is expressly defined for the application as "a subset of the workload", the workload being "the set of identifiable tasks that execute in the processing system" (see: page 8, line 19-page 9, line 3). For each subset of the workload, an impact value is assigned, representing the change in system expiration date that would occur if the workload unit was removed from the system. The term "expiration date" is the date when the server workload is expected to exceed its capacity because of growth in workload. The expiration date may be calculated using life expectancy, capacity space, or other method, as detailed in the Specification on page 11, lines 1-18.

Applicants expressly claim a method (independent Claims 1 and 21 and the claims which depend therefrom), apparatus (Claims 10, 11, 25 and 26) and a program storage device (Claims 12 and 23 and the claims which depend therefrom) for evaluating workload across a processing environment and for altering the workload based on impact value of workload units.

The Yang patent, is directed to a method and apparatus for modeling or profiling a system based on workload in order design the system. A model is generated using "a set of generic system activities" and performance estimates, after which hardware parameters are determined to design a system to handle the activities. Yang uses computer

activity elements ("CAE") "to capture the fundamental activities being performed by [a] respective software application" (see: Col. 4, line 1-8). A set of CAE is calculated for each transaction (see: Col. 26, lines 18-19). A "user focused workload" ("UFW") is collected "using forms or templates ...regarding the application that is available to the user" (Col. 5, lines 50-53). The UFW is translated into a computer activity workload ("CAW") representing the workload data structure that mathematically describes the workload in terms of CAE elements (Col. 5, lines 58-62). Yang models expected application workload for a user and then recommends system requirements for creating a system for that specific user for the user's expected usage of the system. As is clearly taught by Yang, the hardware requirements are determined based on modeled workload. Step 112 in Fig. 1 expressly shows that the final step of Yang's process is to determine hardware parameters. Similarly, step 410 of Fig. 4 shows a final step to "profile hardware system" and step 510 in Fig. 5 shows the final step to "determine possible hardware configurations". Yang is not analyzing an existing processing environment with existing computer systems having assigned workloads and trying to determine how to re-distribute workload among them. Rather, Yang is modeling anticipated workload and designing a system to accommodate that workload. Yang does not determine the

YOR920000461**US1** **-18-**

impact value of existing workload units in order to juggle those workload units among existing computer systems. Yang determines anticipated CAEs and then recommends an appropriate amount of hardware resources to accommodate the number of CAEs workload.

Applicants respectfully assert that the Yang patent neither teaches nor suggests the invention as claimed. With specific reference to the language of the independent claims 1, 10, 12, 21, 23, and 25, Yang does not teach or suggest evaluating workload across an actual processing environment having a plurality of computer systems each having a plurality of assigned existing workload units. Rather, Yang models anticipated requirements for projected applications under user-estimated usage conditions.

The Examiner has, in the **Response to Arguments** section, stated that the recitation of "a processing environment having a plurality of computer systems each having a plurality of assigned workload units" has not been given patentable weight because the language appears in the preamble. Applicants respectfully point out that the same language appears in the body of the claims and that the claimed means and process steps cannot "stand alone" without that language. The claims clearly recite assigning impact values to workload units on each of the plurality of computer systems in a processing environment, calculating

YOR920000461US1 -19-

the impact of moving workload units from one computer system to another computer system, and reassigning the workload among at least two of a plurality of computer systems. Clearly, the body of the claim depends on the existence of a computing environment that comprises a plurality of computer systems and expressly includes steps and means for evaluating and distributing workload among a plurality of computer systems.

The Examiner has further stated, in the **Response to Arguments** section, "that the features upon which applicant relies (i.e., impact values relative to existing workload) are not recited in the rejected claim(s)" citing *In re Van Geuns*. Applicants have amended the language of the independent claims, and of many of the dependent claims, to expressly recite the "existing workload" and "existing workload units". The Amendments are clearly supported by the original Specification (see: e.g., "current location of the workload" page 15, line 25, and "a table of the workload that currently uses system s" page 18, lines 12-13). Applicants further note that the Specification and claims expressly teach that the evaluating includes determining a change in system expiration date. As taught on page 3, lines 18, et seq. the "date when a server's workload exceeds its capacity is called its expiration date" (see also, page

10, lines 14-25). The Specification further teaches that “[t]o provide a quality service, it is necessary to upgrade or offload the server before its expiration date” (page 3, lines 21-22) and “planning dates are chosen in such a way to minimize impact on service” (page 14, lines 14-15). Determining and utilizing the system expiration date as compared to a target planning date for servicing the system with minimal impact on service necessarily involve existing server/resource workload.

The Examiner has concluded that the Applicants’ use of the terms “evaluating”, “assessing” and “planning” does not “amount to existing workload, but rather future performance and potential for workload distribution, thus analogous with Yang’s capacity planning system” (pages 6-7 of the OA). Applicants respectfully disagree and refer the Examiner’s attention to the above-cited passages of the Specification where the use of current workload is clearly described (see also, page 12, lines 16-18 wherein Applicants teach that “[t]he resources consumed by a workload unit are recorded by the processing system and the record of this consumption is transferred as the workload usage history to the administrative processor”). Clearly the present invention is directed to evaluating and distributing existing workload in a currently operating computing environment, and not just in a projected computing system that is under development.

Applicants herein amend the claims to include the language "evaluating and distributing existing workload" to more clearly distinguish the present invention from prior art work related to developing computing systems. Applicants further replace the claim language "assigning impact values" with "calculating impact values" (see: reference numeral 702 of Fig. 7 for support and page 5, line 7-9, the measure/impact value of a workload unit "is calculated in relation to a specific processing system and the other workload assigned to the system"). Finally, Applicants have amended the claims by replacing the term "planning date for servicing" with "service date" (see: original Claims 6 and 15 for support). While the service date is still "planned" insofar as it is a date designated in the future for servicing the computer system, it is nonetheless a target date for actual servicing of a currently working computer system having an existing workload. Comparing a planned service date with a calculated system expiration date (reference numerals 405 and 406 of Fig. 4) would simply not be applicable to a development environment.

In rejecting the steps and means for assigning a plurality of impact values to assigned workload units (Claims 1, 10, 12 and 22, 24 and 26), the Examiner generally cites the teachings from Col. 5, line 1 to Col. 6, line 19. Applicants respectfully assert that the cited passage does

not teach or suggest the assigning of impact values. Yang generates a workload data structure to mathematically describe workload estimates by assigning CAEs in a processing "vacuum", such that an existing processing environment and existing computer system's workload is not taken into account. In contrast, and as specifically taught by the present Specification at page 5, line 7-9, the measure/impact value of a workload unit "is calculated in relation to a specific processing system and the other workload assigned to the system". Where the current claims expressly recite the assignment of impact values relative to the existing workload and system expiration date, Yang uses theoretical values for defining CAEs and the amount of resources required for CAEs (see: the cited passage from Col. 5, line 1 to Col. 6, line 19). While Yang may use prior/historical averages of resource requirements based on data from other systems (Col. 6, lines 9-12), Yang is neither teaching nor suggesting determining impact values for an existing system having an assigned workload.

In rejecting the independent claim features of the steps and means for assessing the impact of moving a workload unit from a donor computer system to a recipient computer system based on impact values (Claims 1, 10 and 12 and Claims 21-26), the Examiner has again cited the passage from Col. 33, lines 30-62. What Yang teaches in the cited

passage is that different possible hardware configurations are considered for a modeled workload. Yang looks at the resources of a hardware configuration and determines whether the resources are sufficient to handle the modeled CAW. If not, Yang looks at a different projected hardware configuration to handle the CAW. If more than one hardware configuration is suitable, Yang may further look at other "desirability" factors (e.g., cost) in selecting a hardware configuration. Yang neither teaches nor suggests assessing the impact of moving some of the CAEs of the CAW to a different hardware configuration.

Yang does not teach or suggest that a system has an expiration date, since Yang does not have a predefined, existing processing environment having a plurality of computer systems each having a plurality of assigned workload units. Rather, Yang models a larger system to accommodate larger workloads, so the concept of system expiration date is simply not applicable. Applicants again note that the Examiner has not cited any teachings in Yang against the "system expiration" language. Rather, the Examiner states "the use and advantages for using such a system is well known to one skilled in the art at the time the invention was made as evidenced by the teachings of Chen".

The Chen patent is directed to a system having parallel database locations for storing data. When a new node (i.e., storage location) is added to the system or when an existing node has excess capacity (see: Abstract), data is moved from one storage node to the new or existing available node. The Examiner has cited Col. 4, lines 56-61 of Chen as providing the teachings missing from the Yang patent, namely "a processing environment each computer system having a plurality of assigned workload units and altering the workload in the processing environment to change expiration dates of at least two of said plurality of computer systems". The cited passage from Col. 4 of Chen merely states that multiple computer systems are coupled to a network, but does not teach or suggest workload units, altering workload, or changing expiration dates of computer systems. The Examiner additionally cites Col. 6, lines 3-32 of Chen. The cited Col. 6 passage describes the data redistribution process of Fig. 6 but states that "selecting the buckets to be redistributed is outside the scope of this present invention" citing a prior article and a patent application. The passage does not teach or suggest computing system workload units, altering workload among a plurality of computing systems, or changing expiration dates of computer systems. Mention is made of selecting buckets of stored data for redistribution "to minimize the impact to

YOR920000461US1 -25-

the performance (*sic*) of the PDB system" and that value "is measured by determining the amount of workload in a bucket and the cost of moving the bucket from node to node". Workload in a bucket of stored data is not defined by Chen; however, it clearly does not refer to processing workload, since Chen is moving buckets of stored data and not actual processing workload. Applicants conclude that the Chen teachings of moving data to a different data storage device is not the same as or suggestive of moving existing processing workload from one computing system to another based on computing system expiration date.

In acknowledging that Yang and Chen fail to teach assigning an impact number representing the number of days that the expiration date of the computer system would be changed with all other workload units remaining the same, the Examiner has newly cited the Miller patent. Miller teaches modifying tasks, themselves (Col. 3, lines 29-37), to require less computing capability. Miller does not teach or suggest assigning impact values to workload units as they relate to system expiration date and then assessing the effect of moving workload units from one computer system to another. Rather, Miller expressly teaches modifying the workload units, themselves, (see: Fig. 2, step 52) so that the workload units require less processing time or fewer resources. The cited passage from Col. 9, lines 3-68 teach

YOR920000461US1 -26-

the initial Miller assessment of task requirements and resource availability, but do not teach or suggest assigning impact values related to system expiration date. Further, since Miller teaches away from moving workload among resources and instead teaches modifying tasks at assigned resources, it cannot be concluded that Miller provides the teachings which are missing from the combination of Yang and Chen.

For a determination of obviousness, the prior art must teach or suggest all of the claim limitations. "All words in a claim must be considered in judging the patentability of that claim against the prior art" (In re Wilson, 424 F. 2d 1382, 1385, 165 U.S.P.Q. 494, 496 (C.C.P.A. 1970)). Since the cited Yang and Chen references fail to teach each and every one of the claim limitations, a *prima facie* case of obviousness has not been established by the Examiner.

With specific reference to the language of the dependent claims, Applicants note that the Examiner has repeated the previous rejections and citations. Applicants, therefore, maintain the contention that the combination of references does not obviate the claims.

With regard to Claims 4 and 13, the Examiner has again cited the passage, and code detailed therein, from Col. 26, line 45-Col. 27, line 5. Yang detailed calculating anticipated reads and writes for one transaction to model YOR920000461**US1** **-27-**

the workload for the transaction. Such is not the same as or suggestive of sorting actual workload units based on assigned impact values.

With regard to the language of Claims 5 and 14, which has been added to the independent claims by this amendment, the Examiner again cited the passage from Col. 25, lines 13-20 which teaches characterizing one transaction in terms of CAE. There is not mention in the passage of altering workload of one or more system. Applicants reiterate that since Yang is "operating" in a theoretical environment, Yang neither teaches nor suggests actual expiration dates of actual systems or altering actual workloads for one or more systems.

With respect to Claims 6 and 15, and Claims 7-9 and 15-18 which depend respectively therefrom, Applicants argue that the cited passage from Col. 33, lines 30-62 does not anticipate or obviate the claimed step and means for comparing the expiration data of each of a plurality of actual computing systems to at least one target planning date for servicing the actual computing systems. What Yang mentions in the cited passage is response times for modeled systems. Again, since Yang is operating in the theoretical environment, real factors such as servicing dates are simply not relevant. Yang does not teach or suggest any consideration of target planning dates for servicing

YOR920000461**US1** **-28-**

systems. Clearly, therefore, Yang does not teach or suggest the language of Claims 6-9 and 15-18.

In rejecting the language of Claims 7 and 16, the Examiner cites Col. 25, lines 13-20 of Yang. The cited passage states that "the transaction mix can be altered by changing the setting for number of transactions, the frequency, or the percentage mix for each transaction type." Those teachings relate to Yang generating a set of CAE for (i.e., modeling) a transaction. There is nothing in the cited passage which teaches or suggests expiration dates, target service dates for system servicing, or altering workloads for at least two actual computing systems.

With regard to the language of Claims 8 and 17, the Examiner again cites the Yang teachings found in Col. 6, lines 9-36. The cited passage states that workload definition information for use in modeling can be collected from actual workload information, that the workload definition is "transferable to other applications and hardware", that "the same workload definition can be used for analysis performed with respect to a plurality of different hardware platforms". Creating a generic workload definition for system modeling does not anticipate the claims which recite creating "From" and "To" lists relative to at least one target service date for servicing of actual computing systems, and reassigning actual workload units

based on assigned impact values reflecting a change in system expiration date if a workload unit is moved from one actual computing system to another. Yang is providing a generic workload definition that can be used to evaluate different candidate system. Such is not the same as or suggestive of evaluating and moving actual workload units in a runtime environment.

The language of Claims 9 and 18, which recites calculating new expiration dates for computing systems on the "From" and "To" lists after reassignment of workload units, has additionally been rejected. The Examiner cites Col. 5, line 1 through Col. 6, line 36 of Yang. The cited teachings all relate to how Yang defines its theoretical workload and CAE. Applicants again argue that the Yang teachings do not anticipate system expiration dates at all, since it is modeling candidate systems, and do not anticipate calculating new expiration dates of actual systems as claimed. Applicants reiterate that Yang is essentially designing a system and opts to change the system design, but does not teach or suggest altering the workload of an existing system.

With regard to Claim 11, Applicants acknowledge that Fig. 7 of Yang illustrates a storage location. However, Yang's provision of a storage location in a "mechanism for profiling a system" is not the same as or suggestive of YOR920000461US1

providing a storage location in the apparatus of Claim 10, including an administrative processor comprising an impact value component for assigning a plurality of impact values, one impact value for each workload unit assigned to each of the plurality of computing systems in a processing environment having a plurality of computer systems each having a plurality of assigned workload units, wherein said assigning of each impact value comprises determining a change in system expiration date should a workload unit be removed from the system; and a processing component for assessing the workload based on said impact values.

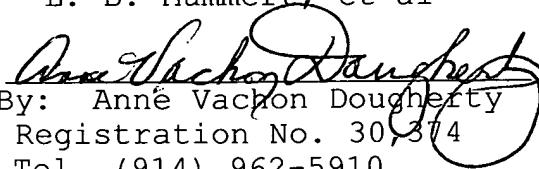
Appellants reiterate that the prior art must teach or suggest all of the claim limitations. "All words in a claim must be considered in judging the patentability of that claim against the prior art" (In re Wilson, 424 F. 2d 1382, 1385, 165 U.S.P.Q. 494, 496 (C.C.P.A. 1970)). Since the cited Yang and Chen references fail to teach each and every one of the claim limitations, a *prima facie* case of obviousness has not been established by the Examiner. Further, obviousness cannot be maintained without some teaching or suggestion of the claim features. The Federal Circuit has stated that when patentability turns on the question of obviousness, the obviousness determination "must be based on objective evidence of record" and that "this precedent has been reinforced in myriad decisions, and

YOR920000461US1 -31-

cannot be dispensed with." (In re Lee, 277 F. 3d 1338, 1343 (Fed. Cir. 2002)). Moreover, the Federal Circuit has stated that "conclusory statements" by an examiner fail to adequately address the factual question of motivation, which is material to patentability and cannot be resolved "on subjective belief and unknown authority" (Id. at 1343-1344). Applicants respectfully assert that the Examiner has not cited actual teachings but has provided conclusory statements about the teachings of Yang, Chen and Miller without basis in actual teachings from those patents.

Based on the foregoing amendments and remarks, Applicants respectfully request entry of the amendments, reconsideration of the amended claim language in light of the remarks, withdrawal of the rejections, and allowance of the claims.

Respectfully submitted,
L. B. Mumment, et al


By: Anne Vachon Dougherty
Registration No. 30,374
Tel. (914) 962-5910